

Complexity and Approximability of Parameterized MAX-CSPs

Holger Dell¹, Eun Jung Kim², Michael Lampis³, Valia Mitsou^{*4}, and Tobias Mömke^{†5}

¹ Saarland University and Cluster of Excellence, Saarbrücken, hdell@mmci.uni-saarland.de

^{2,3} Université Paris Dauphine, eunjungkim78@gmail.com, michail.lampis@dauphine.fr

⁴ SZTAKI, Hungarian Academy of Sciences, Budapest, vmitsou@sztaki.hu

⁵ Saarland University, Saarbrücken, moemke@cs.uni-saarland.de

November 18, 2015

Abstract

We study the optimization version of constraint satisfaction problems (Max-CSPs) in the framework of parameterized complexity; the goal is to compute the maximum fraction of constraints that can be satisfied simultaneously. In standard CSPs, we want to decide whether this fraction equals one. The parameters we investigate are structural measures, such as the treewidth or the clique-width of the variable-constraint incidence graph of the CSP instance.

We consider Max-CSPs with the constraint types AND, OR, PARITY, and MAJORITY, and with various parameters k , and we attempt to fully classify them into the following three cases:

1. The exact optimum can be computed in FPT time.
2. It is W[1]-hard to compute the exact optimum, but there is a randomized FPT approximation scheme (FPT-AS), which computes a $(1 - \epsilon)$ -approximation in time $f(k, \epsilon) \cdot \text{poly}(n)$.
3. There is no FPT-AS unless $\text{FPT} = \text{W}[1]$.

For the corresponding standard CSPs, we establish FPT vs. W[1]-hardness results.

1 Introduction

Constraint Satisfaction Problems (CSPs) play a central role in almost all branches of theoretical computer science. Starting from CNF-SAT, the prototypical NP-complete problem, the computational complexity of CSPs has been widely studied from various

^{*}Supported by ERC Starting Grant PARAMTIGHT (No. 280152)

[†]This research is supported by Deutsche Forschungsgemeinschaft grant BL511/10-1

points of view. In this paper we focus on two aspects of CSP complexity which, though extremely well-investigated, have mostly been considered separately so far in the literature: parameterized complexity and approximability. We study four standard predicates and contribute some of the first results indicating that the point of view of approximability considerably enriches the parameterized complexity landscape of CSPs.

Parameterized CSPs. The vast majority of interesting CSPs are NP-hard [28, 18]. This has motivated the study of such problems from a parameterized complexity point of view, and indeed this topic has attracted considerable attention in the literature [14, 33, 11, 25, 13, 32]. We refer the reader to [27] where an extensive classification of CSP problems for a large range of parameters is given. In this paper we focus on *structurally* parameterized CSPs, that is, we consider CSPs where the parameter is some measure of the structure of the input instance. The central idea behind this approach is to represent the structure of the CSP using a (hyper-)graph and leverage the powerful tools commonly applied to parameterized graph problems (such as tree decompositions) to solve the CSP.

The typical goal of this line of research is to find the most general parameterization of a CSP that still remains fixed-parameter tractable (FPT). To give a concrete example for a very well-known CSP, CNF-SAT is FPT when parameterized by the treewidth of its incidence graph¹ [30] but it is W-hard for more general parameters such as clique-width [23], or even the more restricted modular treewidth [24]. General (boolean) CSP on the other hand, where the description of each constraint is part of the input is known to be a harder problem: it is already W[1]-hard parameterized by the incidence treewidth, but FPT parameterized by the treewidth of the primal graph [31]. Thus, parameterized investigations aim to locate the boundary where a CSP jumps from being FPT to being W-hard. It is of course a natural question how we can deal with the W-hard cases of a CSP once they are identified.

Approximation. CSPs also play a central role in the theory of (polynomial-time) approximation algorithms [34, 19, 4]. In this context we typically consider a CSP as an optimization problem (MAX-CSP) where the goal is to find an assignment to the variables that satisfies as many of the constraints as possible. Unfortunately, essentially all non-trivial CSPs are hard to approximate (APX-hard) from this point of view [6, 18], even those where deciding if an assignment can satisfy all constraints is in P (e.g. 2CNF-SAT or HORN SAT). Thus, research in this area typically focuses on discovering exactly the best approximation ratio that can be achieved in polynomial time. Amazingly, for many natural CSPs this happens to be exactly the ratio achieved by a completely random assignment [17]. This motivates the question of whether we can find natural cases where non-trivial efficient approximations are possible.

Results. In this paper we consider four different types of CSPs where the constraints are respectively OR, AND, PARITY and MAJORITY functions. Our approach follows, for the most part, the standard parameterized complexity script: we consider the input

¹See the next section for a definition of incidence graphs

instance’s incidence graph and try to determine the complexity of the CSP when parameterized by various graph widths. The new ingredient in our approach is that, in addition to trying to determine which parameters make a CSP FPT or W-hard, we also ask if the optimization versions of W-hard cases can be well-approximated. We believe that this is a question of special interest since, as it turns out, there are CSPs for which W-hardness can be (almost) circumvented using approximation, and others which are inapproximable.

More specifically, our results are as follows: for OR constraints, which corresponds to the standard CNF-SAT (MAX-CNF-SAT) problem, we present a new hardness proof establishing that deciding a formula’s satisfiability is W-hard even if parameterized by the incidence graph’s neighborhood diversity. Neighborhood diversity is a parameter much more restricted than modular treewidth (already a restriction of clique-width) [20], for which the strongest previously known W[1]-hardness result was known [24]. We complement this negative result with a strong positive approximation result: there exists a randomized FPT Approximation Scheme (FPT-AS)² for MAX-CNF-SAT parameterized by clique-width, that is, an algorithm which for all $\epsilon > 0$ runs in time $f(k, \epsilon)n^{O(1)}$ and returns an assignment satisfying $(1 - \epsilon)\text{OPT}$ clauses. Thus, even though we establish that solving CNF-SAT exactly is W-hard even for extremely restricted dense graph parameters, MAX-CNF-SAT is well-approximable even in the quite general case of clique-width. To the best of our knowledge, this is the first approximation result of this type for a W-hard MAX-CSP problem.

Recalling that MAX-CNF-SAT is FPT parameterized by the treewidth of the incidence graph, we consider other problems for which the jump from treewidth to clique-width could have interesting complexity consequences. We show that MAX-DNF-SAT and MAX-PARITY, which are FPT parameterized by treewidth, exhibit two wildly different behaviors. On the one hand, the problem of maximizing the largest possible number of satisfied PARITY constraints remains FPT even for dense parameters such as clique-width. On the other hand, by modifying our reduction for CNF-SAT, we are able to show not only that maximizing the number of satisfied AND constraints is W[1]-hard parameterized by neighborhood diversity, but also that this problem cannot even admit an FPT-AS (like MAX-CNF-SAT), unless W[1]=FPT. We recall that PARITY and AND constraints are similar in other aspects: for example, for both we can decide in polynomial time if an assignment satisfying all constraints exists.

Finally, we consider CSPs with MAJORITY constraints, that is, constraints which are satisfied if at least half their literals are true. We give a reduction establishing that this is an interesting case of a natural constraint type for which deciding satisfiability is already W[1]-hard parameterized by treewidth (we actually show W[1]-hardness for the more restricted case of incidence feedback vertex set) and by neighborhood diversity. We complement this negative result with two algorithmic results: first, we show that the corresponding MAX-CSP is FPT parameterized by incidence vertex cover. Then, we use this algorithm as a sub-routine to obtain an FPT-AS for incidence feedback vertex set. Both of these algorithmic results also apply to the more general case of THRESHOLD

²We follow here the standard definition of FPT-AS given in [22].




Figure 1: The structural parameters we study and their relationships. For example, the arrow between cw^* and tw^* means that if the treewidth is bounded, then the clique-width is bounded as well — more precisely, there is a monotone computable function $f : \mathbf{N} \rightarrow \mathbf{N}$ so that $cw^* \leq f(tw^*)$. On the other hand, tw^* and nd^* as well as fvs^* and nd^* cannot be bounded by each other in general.

constraints. We leave it as an interesting open problem to examine if the approximation algorithm for feedback vertex set can be extended to treewidth.

2 Preliminaries

A Boolean CSP ψ is defined as a set $\{C_1, \dots, C_m\}$ of m constraints over a set $X(\psi) = \{x_1, \dots, x_n\}$ of n variables and their negations. Each constraint C_i is regarded as a function of literals (positive or negative appearances of variables) mapped to the set $\{0, 1\}$, where literals can take the values 0 or 1. Furthermore, we define $|C_j|$ to denote the arity of constraint C_j (the number of literals that occur in C) and $|\psi| = m$ the number of constraints in ψ . For simplicity, we also write $l_i \in C_j$ for a literal l_i and a constraint C_j if l_i appears in C_j .

We will be dealing with Boolean constraint satisfaction problems for four well-studied Boolean functions: OR constraints, AND constraints, PARITY (or XOR) constraints and MAJORITY constraints. We say that an assignment $t : X \rightarrow \{0, 1\}$ satisfies a constraint C of type:

- OR, if $\exists l_i \in C, t(l_i) = 1$;
- AND, if $\forall l_i \in C, t(l_i) = 1$;
- PARITY, if it satisfies some equation $\sum_{l_i \in C} t(l_i) = b$ (for $b \in \{0, 1\}$) modulo 2;
- MAJORITY, if at least $\lceil |C|/2 \rceil$ literals in C are set to 1. More generally, we may consider THRESHOLD constraints, where a certain threshold number of literals must be set to true to satisfy the constraint.

Let $\text{occ}(\psi) = \sum_{C \in \psi} |C|$ be the total number of variable occurrences in ψ , that is, the total size of the formula. For a variable x , we write ψ_x for the set of all constraints $C \in \psi$ where x occurs either positive or negative; for the functions we consider without loss of generality, no constraint contains both literals. Thus, the total number of occurrences of a variable x is equal to $|\psi_x|$.

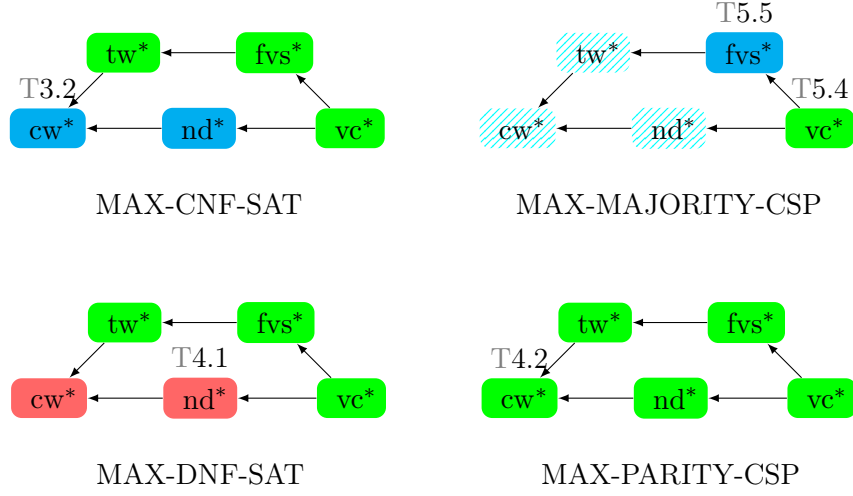


Figure 2: The parameterized complexity status of MAX-CSP problems. The *gray* labels above the boxes indicate the theorem in which we establish the result; previously known results are displayed without reference. *Red* means that the problem is $W[1]$ -hard to compute exactly, and there is no FPT-AS unless $FPT = W[1]$. *Blue* means that the problem is $W[1]$ -hard to compute exactly, and there is an FPT-AS. *Green* means that the problem is FPT to compute exactly. The *blue/white* stripes mean that it's $W[1]$ -hard to compute exactly, and it's open whether there is an FPT-AS.




Figure 3: The parameterized complexity status of CNF-SAT and MAJORITY-CSP. Recall that DNF-SAT and PARITY-CSP are polynomial-time computable. *Red* means that the problem is $W[1]$ -hard and *green* means that the problem is FPT.

We are dealing also with MAX-CSPs, where given a set of constraints ψ , we are interested in finding an assignment to the variables that maximizes the number of satisfied constraints. The natural decision version of this problem is, given a target k , decide whether there exists an assignment that satisfies at least k constraints. Clearly, the problem where we want to decide whether we can satisfy all the constraints is a special case of the above decision problem since we can set $k = m$, but in some cases we consider this simpler decision version, particularly when we want to show hardness.

In the case of OR constraints, the CSP and MAX-CSP problems correspond to the more widely known CNF-SAT and MAX-CNF-SAT problems. In this case we call the constraints *clauses*. When the constraint function is AND, the MAX-CSP problem is called MAX-DNF-SAT. In that case, the constraints are called *terms*. The problem MAX-PARITY is also known as MAX-LIN-2 (satisfy a maximum number of given linear equations modulo 2).

For a CSP ψ , the incidence graph G_ψ^* is defined as the bipartite graph where we construct one vertex v_i for each (unsigned) variable x_i and one vertex u_j for each constraint C_j and connect v_i with u_j if $x_i \in C_j$.

We are interested in parameterizations of the incidence graph $p(G_\psi^*)$ (or simply p^* if G_ψ^* is clear from the context), where p is a structural parameter of G_ψ^* . We are mostly interested in the two most widely studied graph parameters, treewidth tw^* and clique-width cw^* . We refer the reader to standard parameterized complexity textbooks for the definitions, as well as the definitions of standard parameterized complexity terminology used in this paper [8].

3 CNF-SAT and MAX-CNF-SAT

In this section, we consider one of the most fundamental problems in computer science: the satisfiability problem for CNF formulas, which can be viewed as a constraint satisfaction problem where the only allowed constraints are clauses, that is, ORs of literals. An optimal solution for MAX-CNF-SAT can be computed in FPT when parameterized by the treewidth tw^* of the incidence graph [1], and hence CNF-SAT can be solved in the same time. When parameterized by the clique-width cw^* of the incidence graph, all known exact algorithms for CNF-SAT and MAX-CNF-SAT run in XP time [29, 26]. Moreover, we don't expect these problems to be in FPT since they are both W[1]-hard parameterized by cw^* [24].

In Section 3.1, we construct an approximation scheme for MAX-CNF-SAT that runs in FPT time. Intuitively, our algorithm works as follows: given a formula ϕ with 'small' incidence clique-width, we first examine the formula to see if it contains many or few 'large' clauses. If the formula contains relatively few large clauses, then we simply disregard them. We then know that the incidence graph does not contain 'large' bi-cliques, so by a theorem of Gurski and Wanke [15] the remaining formula has small treewidth and we can solve the problem. If on the other hand the original formula contains almost exclusively large clauses, then we observe that we can rely on a random assignment to satisfy almost everything.

The hard part of our algorithm is then how to deal with the general case of formulas that may contain clauses of varying arities, for which we use a combination of the ideas for the two basic cases. In particular, after locating and deleting a negligibly small set of ‘medium’ clauses, we use a counting argument to find a set of variables that appear almost exclusively in large clauses. By setting these variables randomly we satisfy almost all large clauses, and we can then use treewidth to handle the remaining instance.

In Section 3.2, we explore a class of CSP instances that is smaller than the class of bounded incidence clique-width instances; our goal is to understand which incidence graph parameter is responsible for the transition from FPT to W[1]. To this end, we have to look for a graph parameter that is bounded by a function of cw^* (where the problem is hard) but can leave the tw^* unbounded (where it’s FPT). In fact, [24] shows that the problem is W[1]-hard parameterized by the modular treewidth mtw^* of the incidence graph, which lies between cw^* and tw^* .³ We study the incidence *neighborhood diversity* nd^* , which is incomparable to tw^* ; however, mtw^* is bounded when nd^* is. We prove that CNF-SAT remains W[1]-hard parameterized by nd^* .

Definition 3.1. *A graph $G(V, E)$ has neighborhood diversity $\text{nd}(G) = k$ if we can partition V into k sets V_1, \dots, V_k such that, for all $v \in V$ and all $i \in \{1, \dots, k\}$, either v is adjacent to every vertex in V_i or it is adjacent to none of them.*

In other words, $\text{nd}(G) = k$ if V can be partitioned into k modules that are either cliques or independent sets.

Formulas whose incidence graph has neighborhood diversity at most k seem very restrictive: there are only at most k variable- and clause-types, where all variables of the same type belong to the same clauses and all clauses of the same type involve the same variables. This class of formulas is a subset of formulas with $\text{mtw}^* \leq k$ because contracting all modules leaves a graph of order at most k , which trivially has treewidth at most k .

3.1 Approximation Algorithm parameterized by clique-width

Theorem 3.2. *There is a randomized algorithm that, given a CNF formula ψ with n variables, m clauses, and incidence clique-width cw , runs in time $f(\epsilon, \text{cw}) \cdot \text{poly}(n + m)$, and outputs a truth assignment that satisfies at least $(1 - \epsilon) \cdot \text{OPT}$ clauses in expectation.*

We formulate the following basic lemma about probability distributions.

Lemma 3.3. *For all $\epsilon, L > 0$ there is a $c = c(\epsilon, L) > 0$ such that all $c' \geq c$ and all sequences $p_1, \dots, p_{c'} \geq 0$ with $\sum_{i=1}^{c'} p_i \leq 1$ have an index $d \leq c/L$ with the property*

$$p_{[d, L \cdot d]} \doteq \sum_{j=d}^{L \cdot d} p_j < \epsilon.$$

³ A graph of *bounded modular treewidth* is a graph of bounded treewidth after merging modules into a single vertex, where a module is a set of vertices with same neighborhood outside of the set. In fact, the reduction in [24] constructs a formula whose incidence graph has small feedback vertex set after contracting modules.

Proof. Let $\epsilon, L > 0$. We set $c = c(\epsilon, L)$ below. Assume for contradiction that $p_{[d, L \cdot d]} \geq \epsilon$ holds for all $d \in [1, c/L]$. If there are $1/\epsilon + 1$ disjoint intervals $[a_1, L \cdot a_1], \dots, [a_{1/\epsilon+1}, L \cdot a_{1/\epsilon+1}] \subseteq [1, c]$, then we arrive at a contradiction with the fact that the p_i 's are non-negative and sum to at most one. Clearly there exists a constant $c = c(\epsilon, L)$ such that $1/\epsilon + 1$ disjoint intervals of the form $[a, La]$ fit into $[1, c]$. This proves the claim. \square

For an arbitrary given $\epsilon > 0$, we fix $L = \epsilon^{-4}$. We use Lemma 3.3 as follows: For a CNF formula ψ , we define p_i as the fraction of clauses of size i , that is,

$$p_i \doteq \frac{\left| \left\{ C \in \psi \mid |C| = i \right\} \right|}{|\psi|}.$$

Then Lemma 3.3 gives us a number $d \leq c(\epsilon)$ such that the total fraction of clauses whose size is between d and $\epsilon^{-4}d$ is bounded by ϵ . It is now natural to partition all clauses into short, medium, and long clauses. More precisely, we define $\psi = \psi^{<d} \dot{\cup} \psi^{[d, D]} \dot{\cup} \psi^{>D}$ for $D = \epsilon^{-4}d$ as follows:

$$\begin{aligned} \psi^{<d} &\doteq \left\{ C \in \psi \mid |C| < d \right\}, \\ \psi^{[d, D]} &\doteq \left\{ C \in \psi \mid d \leq |C| \leq D \right\}, \text{ and} \\ \psi^{>D} &\doteq \left\{ C \in \psi \mid |C| > D \right\}. \end{aligned}$$

An immediate corollary to Lemma 3.3 is thus that we can choose $d \leq c(\epsilon)$ in such a way that $|\psi^{[d, D]}| \leq \epsilon |\psi|$.

Corollary 3.4. *For all $\epsilon > 0$ there is some $c = c(\epsilon) > 0$ such that all CNF formulas ψ have some $d = d(\epsilon) \in [1, c]$ with $|\psi^{[d, \epsilon^{-4}d]}| \leq \epsilon \cdot |\psi|$.*

If $\psi^{[d, D]} = \emptyset$ holds for $D = \epsilon^{-4}d$ and $d \in [1, c(\epsilon)]$, we say that ψ is ϵ -well separated. We call ψ ϵ' -balanced if, in addition, we have $|\psi^{<d}| \geq \epsilon' m$ and $|\psi^{>D}| \geq \epsilon' m$.

Lemma 3.5. *Let ψ be an ϵ -well separated formula (and thus $V = V(\psi^{<d}) \cup V(\psi^{>D})$).*

Then, for each subset $\hat{\psi} \subseteq \psi^{>D}$ with $|\hat{\psi}| > \epsilon^2 m$, there is a variable y such that $|\psi_y^{<d}| \leq \epsilon^2 |\hat{\psi}_y|$.

That is, for every set $\hat{\psi}$ that contains a significant fraction of long clauses, there is a variable that occurs $|\hat{\psi}_y|$ times in $\hat{\psi}$, but only at most an ϵ^2 -fraction of that in the short clauses.

Proof. Let $\hat{\psi} \subseteq \psi^{>D}$ with $|\hat{\psi}| > \epsilon^2 m$. Note that the total number of literal occurrences in $\hat{\psi}$ is $\text{occ}(\hat{\psi}) > D \cdot \epsilon^2 \cdot m = \epsilon^{-2} dm$. In contrast, $\text{occ}(\psi^{<d}) < dm$. Now suppose that there was no variable y with the claimed properties, that is, suppose that every variable y satisfies $|\psi_y^{<d}| > \epsilon^2 |\hat{\psi}_y|$. Then the total number of variable occurrences in $\psi^{<d}$ can be bounded from below as follows:

$$\text{occ}(\psi^{<d}) = \sum_y |\psi_y^{<d}| > \sum_y \epsilon^2 |\hat{\psi}_y| = \epsilon^2 \text{occ}(\hat{\psi}) > d \cdot m.$$

This yields a contradiction and thus proves the claim. \square

Proof of Theorem 3.2. The algorithm A works as follows. Let $\epsilon' = \epsilon^2$, and we assume w.l.o.g. that $\epsilon < 1/8$. Given a CNF formula ϕ , we compute an ϵ' -well separated formula ψ by dropping all clauses in $\phi^{[d,D]}$. Corollary 3.4 guarantees that the fraction of deleted clauses is bounded by ϵ' . If ψ is not $\epsilon/2$ -balanced, we discard the smaller side (with fewer clauses) and only handle the larger one: If $\psi^{<d}$ is the larger side, we compute an optimal assignment for $\psi^{<d}$ in FPT time, by using the result of Gurski and Wanke [15]. This way the total number of unsatisfied clauses is at most $\epsilon m/2$, and together with the unsatisfied clauses due to applying Corollary 3.4, the total number of unsatisfied clauses is smaller than ϵm . Since $\text{OPT} > m/2$, we get the approximation guarantee.

If $\psi^{>D}$ is the larger side, we use a random assignment. This way, at most $\epsilon m/2$ clauses from $\psi^{<d}$ are violated, and in expectation at most a 2^{-D} fraction of clauses from $\psi^{>D}$ are violated. Since 2^{-D} is smaller than $\epsilon/4$, we conclude that – together with unsatisfied clauses due to applying Corollary 3.4 – at least $(1 - \epsilon)m$ clauses are satisfied in expectation.

This finishes the analysis of unbalanced formulas, and in the remaining proof we may assume that ψ is $\epsilon/2$ -balanced. To handle this case, we determine a set of variables Y such that

- there are at most $\epsilon m/4$ short clauses with variables from Y and
- there are at most $\epsilon^2 m$ long clauses that contain $\leq 1/\epsilon$ variables from Y .

Before we construct Y , let us verify that the properties of Y imply the correctness of the theorem. Our algorithm computes a satisfying assignment of the short clauses without variables from Y , again using the result of Gurski and Wanke [15]. Afterwards it assigns values uniformly at random to the variables in Y .

There are at most $\epsilon' m = \epsilon^2 m$ unsatisfied clauses due to applying Corollary 3.4, $\epsilon m/4$ short clauses clauses that we did not consider when satisfying clauses from $\psi^{<d}$, and $\epsilon^2 m$ clauses from $\psi^{>D}$ that we did not consider in the random assignment. Additionally, in expectation there are less than $2^{-1/\epsilon} m$ clauses left unsatisfied from the remaining $|\psi^{>D}| - \epsilon^2 m$ clauses from $\psi^{>D}$. Since we assumed that $\epsilon < 1/8$, the theorem follows.

To construct the set Y , we iteratively apply Lemma 3.5 with the parameter $\epsilon/4$. Initially, we set $\hat{\psi} = \psi^{>D}$. In each iteration, we identify a variable y according to the lemma and add the variable to Y . In the subsequent iterations, we mark y to be inactive and handle it as if it was not contained in any clause. Whenever we identify a clause C that has at least $1/\epsilon$ inactive variables (i.e., variables from Y), we remove C from $\hat{\psi}$. We continue this process until $|\hat{\psi}| \leq \epsilon^2$. Note that applying Lemma 3.5 for $\epsilon/4$ but having an ϵ' -well separated formula ensures that at all times, all clauses in $\hat{\psi}$ have sufficiently many literals to apply Lemma 3.5. Therefore the process terminates and the generated set Y has the aimed-for properties since $|Y| \leq m/\epsilon$. \square

3.2 Hardness parameterized by neighborhood diversity

A *constraint* on r variables is a relation $R \subseteq \{0, 1\}^r$. We define the unary constraints $U_0 = \{0\}$ and $U_1 = \{1\}$, which corresponds to clauses $(\neg x)$ and (x) , respectively. We

define the equality $=$ and disequality \neq constraints on two groups of Boolean variables $x = x_1 \dots x_n$ and $y = y_1 \dots y_n$ in infix notation in the usual way: For an assignment a to the x - and y -variables, we say that $a \models x = y$ if and only if, for all $i \in [n]$, we have $a(x_i) = a(y_i)$, that is, the assignment sets x_i to the same value as y_i ; as usual, $x \neq y$ is interpreted as the negation of $x = y$.

Lemma 3.6. *CNF-SAT parameterized by nd^* is $\text{W}[1]$ -hard, where nd^* is the neighborhood diversity of the input's incidence graph.*

Proof. We devise an FPT-reduction from k -Multicolored Clique to CNF-SAT. Given a k -partite graph G , whose parts V_1, \dots, V_k all have the same size n , we construct k groups of variables x_1, \dots, x_k , which together are supposed to represent a k -clique in G , should one exist. Each group x_i consists of $\log n$ Boolean variables and represents the supposed clique's vertex in the part V_i . Without loss of generality, we assume that $\log n$ is an integer.

Starting from the empty CNF formula, we construct a formula ϕ on the x -variables as follows. First choose, for each $i \in [k]$, an arbitrary bijection $\text{bin}_i : V_i \rightarrow \{0, 1\}^{\log n}$ that maps any vertex $u \in V_i$ to its binary representation $\text{bin}(u)$; for convenience, we drop the index i . For each $i, j \in [k]$ with $i < j$, and for each non-edge $(u, v) \notin E(V_i, V_j)$ between V_i and V_j , we add the following constraint $C_{i,j,u,v}$ to ϕ :

$$x_i x_j \neq \text{bin}(u) \text{bin}(v).$$

Clearly, this constraint excludes exactly one of the $2^{2 \log n}$ possible assignments to $x_i x_j$, and so it can be written as an OR of literals of the x -variables. In the end, ϕ is a CNF formula with $|E(G)|$ clauses.

For the completeness of the reduction, let $v_i \in V_i$ for all $i \in [k]$ be such that v_1, \dots, v_k is a clique in G . For all $i \in [k]$, set $x_i = \text{bin}(v_i)$. This assignment satisfies all constraints: for all $(u, v) \notin E(V_i, V_j)$, we have that $\text{bin}(v_i) \text{bin}(v_j) \neq \text{bin}(u) \text{bin}(v)$ because (v_i, v_j) is an edge and (u, v) is not, and bin is a bijection.

For the soundness of the reduction, let $a_1, \dots, a_k \in \{0, 1\}^{\log n}$ be a satisfying assignment of ϕ . For each $i \in [k]$, let v_i be the unique vertex in V_i for which $\text{bin}(v_i) = a_i$. Let $i, j \in [k]$ with $i < j$. Since the assignment satisfies all constraints of ϕ , it must be the case that (v_i, v_j) is an edge in G . For if it was a non-edge, its corresponding constraint in ϕ would have excluded the assignment $a_i a_j$ for $x_i x_j$. Hence v_1, \dots, v_k is a clique in G .

It remains to argue that the neighborhood diversity of the incidence graph of ϕ is at most $k + \binom{k}{2}$. The modules of the incidence graph are the variable group x_h for each $h \in [k]$ and the clause group $\{C_{i,j,u,v}\}$ for each $i, j \in [k]$ with $i < j$. Indeed, the incidence graph between x_h and $C_{i,j,*}$ is a bipartite clique if $h \in \{i, j\}$, and otherwise it is an independent set.

We constructed an FPT-reduction from the $\text{W}[1]$ -complete problem Multicolored Clique to CNF-SAT parameterized by nd^* , which finishes the proof of the theorem. \square

4 From Treewidth to Clique-width

In the previous section, we have seen that MAX-CNF-SAT is fixed-parameter tractable when parameterized by tw^* , which is a sparse graph parameter, and it is W[1]-hard to compute exactly and has an FPT-AS when parameterized by nd^* , which is a dense graph parameter. In this section we observe that the transition from sparse to dense parameters has different effects on the complexity of MAX-CSP, depending on which types of constraints are allowed.

By modifying our reduction for CNF-SAT we show that MAX-DNF-SAT, the problem of maximizing the number of satisfied AND constraints is W[1]-hard parameterized nd^* . Furthermore, because the maximum number of constraints that could be satisfied in our reduction is also bounded by some function of the parameter, we show that the problem does not have an FPT-AS unless $\text{FPT} = \text{W}[1]$. Thus, while MAX-DNF-SAT is FPT parameterized by tw^* , it does not even appear to have an FPT approximation scheme when parameterized by nd^* .

Theorem 4.1. *Suppose that there exists an FPT-AS which, given $\epsilon > 0$ and an instance ϕ of MAX-DNF-SAT, computes a $(1 - \epsilon)$ -approximate solution and runs in time $f(\text{nd}^*, \epsilon) \cdot \text{poly}(n)$, where nd^* is the neighborhood diversity of the incidence graph of ϕ . Then $\text{FPT} = \text{W}[1]$.*

Proof. We devise an FPT-reduction from k -Multicolored Clique that is similar to the one in our proof of Theorem 3.6.

Given a k -partite graph G whose parts V_1, \dots, V_k have size n each, we construct k groups x_1, \dots, x_k of $\log n$ variables each, and $\binom{k}{2}$ groups of AND constraints $C_{i,j,u,v}$ for each integers $i, j \in [k]$ with $i < j$ and edge $(u, v) \in E(V_i, V_j)$ between V_i and V_j :

$$x_i x_j = \text{bin}(u) \text{bin}(v)$$

Here, $\text{bin}(u) \in \{0, 1\}^{\log n}$ is some binary representation of $u \in V_i$. Note that the constraint $C_{i,j,u,v}$ is satisfied by exactly one of the $2^{2 \log n}$ assignments to the variables $x_i x_j$, and so this constraint can be written as an AND of literals of these variables. Apart from producing ϕ , the reduction also sets the approximation parameter $\epsilon = k^{-2}$ so that $(1 - \epsilon) \binom{k}{2} > \binom{k}{2} - 1$ holds. The neighborhood diversity of the incidence graph of ϕ is at most $k + \binom{k}{2}$ by a completely analogous argument as in the proof of Theorem 3.6.

We now prove that an FPT-AS for MAX-DNF-SAT would allow us to distinguish whether G has a k -clique or not. For the completeness of the reduction, let G have a k -clique v_1, \dots, v_k . Then the assignment $x_1 \dots x_k = \text{bin}(v_1) \dots \text{bin}(v_k)$ satisfies the $\binom{k}{2}$ constraints C_{i,j,v_i,v_j} for each $i, j \in [k]$ with $i < j$. Thus the assumed FPT-AS will return a solution that satisfies at least $(1 - \epsilon) \binom{k}{2} > \binom{k}{2} - 1$ constraints. Since the number of satisfied constraints is an integer, it must thus be at least $\binom{k}{2}$ (and in fact is equal to $\binom{k}{2}$ in this case). For the soundness, if G has no k -clique, then at most $\binom{k}{2} - 1$ constraints can be simultaneously satisfied, and so the assumed FPT-AS can not return a solution whose value is larger than that.

Finally, note that the overall algorithm above solves the Multicolor Clique problem in time $f(\text{nd}^*, \epsilon) \text{poly}(n) \leq g(k) \text{poly}(n)$, which is FPT; thus an FPT-AS for MAX-DNF-SAT would imply that $\text{FPT} = \text{W}[1]$. \square

When parameterized by tw^* , MAX-CNF-SAT and MAX-DNF-SAT are both FPT, and when parameterized by a dense graph parameter, such as nd^* , the former problem is hard but approximable while the latter problem is hard even to approximate. We next consider natural constraint types where the corresponding CSPs stay FPT both for sparse as well as dense incidence graph parameters. MAX-PARITY wants to find an assignment that satisfies the maximum number of linear equations modulo two. While deciding whether there is an assignment that satisfies all equations is in P (by Gauss elimination), the maximization version is a typical APX-hard problem [17]. Here we show that computing the optimal solution of MAX-PARITY is FPT, regardless of whether the parameter is the treewidth or the clique-width of the incidence graph. Our intuition for why MAX-PARITY appears to be so much easier than CNF-SAT is that negations are (almost) irrelevant, and so the incidence graph seems to capture most of the structure relevant to the complexity of the CSP instance.

Theorem 4.2. *Given an instance ϕ for MAX-PARITY, we can find an optimal solution in time $f(\text{cw}^*)|\phi|^{O(1)}$, where cw^* is the clique-width of the incidence graph of ϕ .*

Proof. We rely on the meta-theorem of [5] that all problem which can be expressed in the optimization version of CMSO_1 can be solved exactly in linear time. In the logic CMSO_1 , we are allowed to express problems via first-order formulas with additional second-order quantifiers that are only allowed to quantify over subsets of the universe of the input structure, and with counting constraints that stipulate the cardinalities of these sets modulo a constant number.

To express MAX-PARITY in CMSO_1 , observe that when given a linear equation $\sum_i l_i = b$ over $\text{GF}(2)$, where $b \in \{0, 1\}$ and the l_i are literals (either x_i or $1 - x_i = 1 + x_i$), we may view it equivalently as a constraint of the form $\sum_i x_i = b'$, where all the x_i are variables and $b' = b$ if and only if the original constraint contains an even number of negated literals on the left hand side.

Having performed the above pre-processing we can now express our problem in CMSO_1 . The structure we construct is a bipartite directed graph G with the bipartition $L \dot{\cup} R$, which is represented by an edge relation E . For each linear equation $\sum_i x_i = b$, we introduce a vertex ℓ in L , for each variable we introduce a vertex x_i in R , and we set $E(\ell, x_i)$ if and only if x_i appears in equation ℓ . Moreover, we have the unit constraints U_0 and U_1 , and we set $U_b(\ell)$ if and only if b is the right-hand side of ℓ .

The CMSO_1 formula that we construct is looking for the largest set $S \subseteq L$ of constraint-vertices such that there exists a set of variables $X \subseteq R$ which satisfies the following: every vertex $\ell \in S$ where $U_b(\ell)$ holds has a number of neighbors in X that is equal to b modulo two, for every $b \in \{0, 1\}$. By construction, the maximum such set S corresponds to the maximum set of linear equations that can be satisfied simultaneously by an assignment represented by X . \square

5 Majority and Threshold CSPs

In this section we deal with CSPs where each constraint is a MAJORITY or a THRESHOLD constraint. In this problem, each constraint is supplied with an integer value t (the threshold) and it is satisfied if and only if at least t of its literals are set to true. MAJORITY is the special case of this predicate where t is always equal to half the arity of each constraint. We denote the resulting satisfiability problem with MAJORITY and THRESHOLD, and the resulting MAX-CSPs with MAX-MAJORITY and MAX-THRESHOLD.

MAJORITY and THRESHOLD constraints are of course some of the most natural and well-studied predicates in many contexts: for example, MAX-CSP for such constraints contains the complexity of finding an assignment that satisfies as many inequalities as possible in a 0-1 Integer Linear Program whose coefficients are in $\{-1, 0, 1\}$. This problem, sometimes called MAXIMUM FEASIBLE SUBSET has been well-studied in the literature [9, 3, 2]. MAJORITY constraints also play a central role in learning theory [10, 16] and in hardness of approximation [7].

5.1 Hardness of exact algorithms

We consider the problem whether a CSP with THRESHOLD constraints is satisfiable. This problem is NP-complete. We parameterize the problem by the size fvs^* of the smallest feedback vertex set, or by the neighborhood diversity nd^* of the instance's incidence graph. As we will see, these parameterized problems turn out to be hard, even for the special case of MAJORITY constraints. Thus, neither dense nor sparse incidence graph parameters appear to put the problem in FPT.

In order to ease notation in the upcoming proofs, we note that THRESHOLD-constraints are quite expressive. For example, they can express clauses $\ell_1 \vee \dots \vee \ell_d$ since this is the same as requiring that at least one of the d literals is true. Similarly, stipulating that at least d literals be true is the same as a term $\ell_1 \wedge \dots \wedge \ell_d$. Finally, we can stipulate AT-MOST-ONE(ℓ_1, \dots, ℓ_d), that is, that at most one of the literals is set to true, by using the THRESHOLD-constraint that at least $d - 1$ of the literals $\neg \ell_1, \dots, \neg \ell_d$ are true.

Theorem 5.1. *THRESHOLD parameterized by fvs^* is W[1]-hard.*

Proof. We reduce from the k -Multicolored Clique problem. Let $G(V_1, \dots, V_k, E)$ be a k -partite graph with $|V_i| = n$ for each $i \in \{1, \dots, k\}$, and we regard the vertices of V_i to be identified with integers from 1 to n . Let E_{ij} be the edge set between V_i and V_j for every $1 \leq i < j \leq k$. We construct the output formula ϕ of THRESHOLD using following gadgets for every i and every pair $1 \leq i < j \leq k$:

- *Vertex selection gadget for V_i :* For each vertex $\ell \in V_i$, we create a sequence of ℓ variables P_ℓ for ϕ and name the first and the last variable in the sequence p_ℓ and p'_ℓ respectively (in particular, if $\ell = 1$, we have $p_1 = p'_1$). For every two consecutive variables $y, z \in P_\ell$, we add an OR-constraint $C = (y \vee \neg z)$. These constraints

guarantee that, in any satisfying assignment and for every ℓ , if p'_ℓ is set to true, then all variables in P_ℓ , including p_ℓ , are set to true as well.

Given the n variable sets P_1, \dots, P_n , we add two constraints X_i and Y_i as follows:

$$X_i = \text{AT-MOST-ONE}(p_\ell : \ell \in V_i) \quad \text{and} \quad Y_i = \text{AT-LEAST-ONE}(p'_\ell : \ell \in V_i).$$

The constraint Y_i enforces at least one p'_ℓ is set to true, which propagates through all variables in P_ℓ to p_ℓ . Conversely, the constraint X_i requires that at most one p_ℓ be set to true. Hence, any satisfying assignment will have exactly one $\ell \in V_i$ for which p_ℓ is set to true; moreover, all variables in the set P_ℓ are set to true, and all the variables in $P_{\ell'}$ for $\ell' \in V_i \setminus \{\ell\}$ are set to false.

- *Edge selection gadget for E_{ij}* : The gadget for E_{ij} is similar to that for V_i . For each edge $e = (k, \ell) \in E_{ij}$ with $k \in V_i$ and $\ell \in V_j$, we create a sequence of $n + 1 - \ell$ variables Q_e and name the first and the last variable in the sequence q_e and q'_e respectively. For every two consecutive variables $y, z \in Q_e$, we add an OR-constraint $C = (y \vee \neg z)$. Given the variable sets Q_e for each $e \in E_{ij}$, we add two constraints X_{ij} and Y_{ij} as follows:

$$X_{ij} = \text{AT-MOST-ONE}(q_e : e \in E_{ij}) \quad \text{and} \quad Y_{ij} = \text{AT-LEAST-ONE}(q'_e : e \in E_{ij}).$$

By the same argument as for the vertex selection gadget, we have exactly one variable q_e set to true, in which case all variables of Q_e are set to true and all variables in $Q_{e'}$ are set to false for all $e' \neq e$.

- *Incidence verification gadget between V_i and E_{ij}* : For every edge $e = k\ell \in E_{ij}$ with $k \in V_i$, we create an OR-constraint $C_{ke} = (p_k \vee \neg q_e)$. This guarantees that any satisfying assignment that sets q_e to true also sets p_k to true.
- *Incidence verification gadget between V_j and E_{ij}* : We add two constraints C_{ij} , C'_{ij} and their t -values as

$$C_{ij} = \text{AT-LEAST}_{n+1} \left(\bigcup_{e \in E_{ij}} Q_e \cup \bigcup_{\ell \in V_j} P_\ell \right), \text{ and}$$

$$C'_{ij} = \text{AT-MOST}_{n+1} \left(\bigcup_{e \in E_{ij}} Q_e \cup \bigcup_{\ell \in V_j} P_\ell \right).$$

Any satisfying assignment will thus set exactly $n + 1$ variables to true among all the variables in the Q_e sets for $e \in E_{ij}$ and the P_ℓ sets for $\ell \in V_j$. In particular, there must be a natural number ℓ such that ℓ of the true variables are in the P -sets, and $n + 1 - \ell$ of the true variables are in the Q -sets. By the constraints on the P -variables, the set of true variables is equal to exactly one of the sets $P_{\ell'}$. Since $|P_{\ell'}| = \ell'$ for all $\ell' \in \{1, \dots, n\}$, we have that $\ell > 0$ and that the selected set is exactly P_ℓ . Moreover, exactly one of the sets Q_e is selected due to the constraints

on the Q_e -variables. Since $n + 1 = |Q_e| + |P_\ell| = |Q_e| + \ell$ holds if and only if $e = k\ell$ for some $k \in V_i$, the selected set Q_e must correspond to an edge e incident to $\ell \in V_j$.

For the completeness of the reduction, let S be a k -clique of G . Then we set all ℓ variables of P_ℓ from the vertex selection gadget for V_i to true if $\ell \in V_i$ is chosen for the k -clique S . Also we set all $n + 1 - \ell$ variables of Q_e , where $e = (k, \ell)$, from the edge selection gadget for E_{ij} to true if $k \in V_i$ and $\ell \in V_j$ are chosen for S . It is not difficult to check that this assignment satisfies ϕ .

Conversely, suppose ϕ has a satisfying assignment. From the above argument, for each V_i there is exactly one variable set to true among all p_ℓ 's for $1 \leq \ell \leq n$. Likewise, for each E_{ij} , exactly one variable q_e is *selected*. From the property of the incidence verification gadget between V_i and E_{ij} , whenever q_e from the edge selection gadget for E_{ij} is *selected*, p_k such that k is the vertex in V_i incident with e in G must be set to true. Lastly, by the property of incidence verification gadget between V_j and E_{ij} , that q_e with $e = (k, \ell)$ is selected implies that exactly ℓ variables are set to true in the vertex selection gadget for V_j . From the property of the vertex selection gadget, this means that P_ℓ is selected and thus p_ℓ is set to true. Hence, for every $e = (k, \ell) \in E_{ij}$, whenever q_e is set to true, both p_k and p_ℓ are set to true. By selecting the vertices of G corresponding to $k \in V_i$ such that p_k is set to true, we find a k -clique of G .

Consider the set of constraints

$$F = \{X_i, Y_i : 1 \leq i \leq k\} \cup \{X_{ij}, Y_{ij}, C_{ij}, C'_{ij} : 1 \leq i < j \leq k\}.$$

Consider the graph obtained by deleting F from the incidence graph of ϕ . Each component involves a variable set P_k , for some vertex $k \in V_i$, and Q_e for all edges $e = (k, \ell)$ such that $\ell \in V_j$ for some $j > i$. It is easy to see that each component is in fact a graph obtained by subdividing edges of a star. Hence the incidence graph of ϕ has a feedback vertex set of size $O(k^2)$. This completes the proof. \square

Theorem 5.2. MAJORITY parameterized by fvs^* is $\text{W}[1]$ -hard, where fvs^* is the minimal size of a feedback vertex set of the instance's incidence graph.

Proof. We devise an FPT-reduction from THRESHOLD to MAJORITY that keeps fvs^* the same. The result then follows from Theorem 5.1.

The reduction's input is a CSP ϕ with threshold constraints, and the output is a CSP ϕ' with majority constraints. We transform threshold into majority constraints by adding fresh variables, some of which are forced to be false.

For each constraint C that has the threshold $t(C) = \lceil |C|/2 \rceil + d$ for some integer d , we add d fresh dummy variables y_1, \dots, y_d to obtain a constraint C' whose threshold we set to $t(C') = \lceil |C'|/2 \rceil$; thus, C' is a majority constraint, and we include it into ϕ' . If $d > 0$, we additionally add the constraint $(\neg y_i)$ for each $i \in [d]$ to force the dummy variables to be false.

For the correctness of the reduction, note that any satisfying assignment of ϕ' sets any dummy variable y to false if ϕ' contains the constraint $(\neg y)$, and it otherwise sets y to

true without loss of generality since y occurs exactly once as a positive literal. Thus, any assignment to the old variables satisfies a constraint C if and only if setting the y -variables as just specified sets at least half of the literals in C' to true. Thus ϕ is satisfiable if and only if ϕ' is.

Adding new variables to a constraint corresponds to adding leaves to the corresponding constraint vertex in the incidence graph, which does not create any new cycles. Adding a single unit clause for a y -variable also does not create new cycles as this just adds a leaf to some leaf. Thus $\text{fvs}^*(\phi) = \text{fvs}^*(\phi')$, and our reduction has the claimed properties. \square

Theorem 5.3. THRESHOLD and MAJORITY parameterized by the incidence neighborhood diversity nd^* are $\text{W}[1]$ -hard.

Proof. First note that OR-constraints are special cases of THRESHOLD-constraints, and so the hardness of THRESHOLD follows Lemma 3.6, where we establish the hardness of CNF-SAT. Next we reduce from CNF-SAT to MAJORITY while preserving the neighborhood diversity. For this, let φ be a CNF-formula and consider its incidence graph with neighborhood diversity k . Let G_1, \dots, G_t the $t \leq k$ modules corresponding to constraints of φ . That is, every constraint in G_i depends on the same set of vertices, and in particular, every constraint in G_i has the same arity a_i . Now for each i , we add a set Z_i of $a_i - 1$ fresh dummy variables as well as the constraints (z) forcing all $z \in Z_i$ to true. We convert each constraint of G_i to a MAJORITY-constraint by adding the variable set Z_i to every constraint in G_i . Since the $z \in Z_i$ are forced to true, each constraint always contains $a_i - 1$ variables set to true among its $2a_i - 1$ variables. Thus the new MAJORITY-constraint is satisfied if and only if the old OR-constraint was. The reduction adds at most k groups Z_i of variables, which form modules in the incidence graph, and so the incidence neighborhood diversity of the output instance is at most $2k$. \square

5.2 Exact Algorithm parameterized by vertex cover

Motivated by the negative result of Theorem 5.2 we now investigate the complexity of MAJORITY for more restricted parameters. The first parameter we consider is the vertex cover of the incidence graph. This is a natural, though quite restrictive, parameter which is often considered for problems which are W -hard for treewidth.

Theorem 5.4. MAX-THRESHOLD parameterized by the incidence vertex cover vc^* is FPT.

Proof. Given a CSP ϕ with THRESHOLD-constraints over a variable set X , and a size- k vertex cover S of the incidence graph, we define $S_X, S_\phi \subseteq S$ to be the vertices of S corresponding to variables and constraints, respectively. The algorithm starts by branching into $2^{|S_X|} \leq 2^k$ cases, corresponding to truth assignments σ on the variables in S_X . Since S is a vertex cover, all variables of constraints $C \in \phi \setminus S_\phi$ are contained in S_X , and so we can compute how many of these constraints are satisfied by σ ; let's call this number N_σ . After fixing the variables in S_X and removing the constraints in $\phi \setminus S_\phi$, what remains is a CSP ϕ' with at most k constraints. Let N'_σ be the maximum number of constraints that can be simultaneously satisfied in ϕ' . Then $\max_\sigma N_\sigma + N'_\sigma$ is

the maximum number of constraints of ϕ that can be simultaneously satisfied. Thus it remains to compute N'_σ in FPT-time.

Let ϕ' be a MAX-CSP with k constraints. We reduce it to 2^k instances ϕ'' of the standard CSP-version of THRESHOLD, by guessing for each constraint whether or not it is satisfied by an optimal solution. The size of the largest subset of constraints that can be simultaneously satisfied is then precisely the optimal value of the MAX-CSP-instance.

Finally, to solve THRESHOLD, we further reduce ϕ'' to an integer linear program (ILP) with 3^k variables. This suffices since solving ILPs is FPT when parameterized by the number of variables [21]. Let k be the number of constraints of ϕ'' . We associate with each variable x of ϕ'' the vector $v(x) \in \{-1, 0, 1\}^k$ with

$$v(x)_i = \begin{cases} 1 & \text{if } x \text{ occurs positively in the } i\text{-th constraint,} \\ -1 & \text{if } x \text{ occurs negatively in the } i\text{-th constraint, and} \\ 0 & \text{if } x \text{ does not occur in the } i\text{-th constraint.} \end{cases}$$

For each vector $v \in \{-1, 0, +1\}^k$, we add a variable ℓ_v to the ILP and the constraints $0 \leq \ell_v \leq B_v$, where $B_v \in \mathbf{N}$ is the number of variables x with $v(x) = v$. The variable ℓ_v is supposed to indicate how many of the variables of type v are set to true. Finally, we translate each constraint $C_i \in \phi''$ to a linear equation. Let T_i be the threshold of the constraint, and let n_i be the number of variables that occur as negative literals in C_i . Then we add the following inequality to the ILP:

$$\left(\sum_{\substack{v \in \{-1, 0, 1\}^k \\ v_i = 1}} \ell_v \right) - \left(n_i - \sum_{\substack{v \in \{-1, 0, 1\}^k \\ v_i = -1}} \ell_v \right) \geq T_i. \quad (1)$$

To prove the completeness of this reduction, let ϕ'' have a satisfying assignment σ . Then we set ℓ_v to be the number of variables x with $v(x) = v$ such that $\sigma(x) = 1$. This satisfies $0 \leq \ell_v \leq B_v$. All other linear constraints are generated from some constraint C_i of ϕ'' , which is satisfied by σ . The first term in the difference of (1) is exactly the number of variables x that are set to true under σ and that occur positively in C_i , and the second term is the number of variables x that are set to false and that occur negatively in C_i ; thus, the left-hand side is the number of literals set to true and the right-hand side is the threshold of C_i , and so (1) holds.

For the soundness, assume there is a solution $(\ell_v)_v$ for the ILP. Then we construct an assignment σ as follows: For each v , arbitrarily select ℓ_v of the B_v variables with $v(x) = v$ to true, and set the others to false. Then for each constraint C_i of ϕ'' , the linear constraint (1) guarantees that C_i is satisfied. This finishes the correctness proof of the final reduction; overall, we solve MAX-THRESHOLD in FPT-time when parameterized by vc^* . \square

5.3 Approximation Algorithm parameterized by feedback vertex set

The results of Theorem 5.2 naturally pose the following question: can we evade the W-hardness of MAJORITY by designing an FPT-AS for the problem? In this section,

though we do not resolve this question, we give some first positive indication that this may be possible. We consider THRESHOLD parameterized by the incidence graph's feedback vertex set (that is, the number of vertices that need to be deleted to make the graph acyclic). This is a natural, well-studied parameter that generalizes vertex cover but is a restriction of treewidth. It is also connected to the concept of back-door sets to acyclicity, which is well-studied in the parameterized CSP literature [23, 12].

Observe that approximating this CSP is non-trivial, since MAX-MAJORITY with constraints of arity two already generalizes MAX-2SAT, and is hence APX-hard. On the other hand, MAX-MAJORITY can easily be 2-approximated by considering any assignment and its negation. Hence, the natural goal here is an approximation ratio of $(1 - \epsilon)$. Using Corollary 5.4 as a sub-routine we achieve this with an FPT-AS.

Theorem 5.5. *There exists an FPT-AS which, given $\epsilon > 0$ and an instance ϕ of MAX-THRESHOLD, computes a $(1 - \epsilon)$ -approximate solution and runs in time $f(\text{fvs}^*, \epsilon) \cdot \text{poly}(n)$, where fvs^* is the size of the smallest feedback vertex set of the incidence graph of ϕ .*

Proof. We consider two cases:

- If $|\phi| \leq 2k/\epsilon$, then the number of constraints is bounded by a function of k and ϵ , and we can use the FPT-algorithm from Theorem 5.4.
- If $|\phi| > 2k/\epsilon$, then with a similar argument as in the proof of Theorem 5.4 we can consider that the $\text{fvs}(G_\phi^*)$ contains only constraint vertices (that is, we guess the assignment of variables in the feedback vertex set). We now proceed by simply deleting these constraints from the instance. Since the resulting instance is acyclic we can find the optimal solution in the remaining constraints of ϕ in polynomial time. Call the produced solution SOL and the optimal solution OPT. Then we have: $\text{SOL} \geq \text{OPT} - k$. Now, observe that $\text{OPT} \geq |\phi|/2$ because if an assignment satisfies less than half of the constraints, its negation satisfies the rest. Therefore $\text{OPT} > k/\epsilon$ which gives $\frac{\text{OPT}-k}{\text{OPT}} > 1 - \epsilon$.

□

References

- [1] M. Alekhovich and A. A. Razborov. Satisfiability, branch-width and Tseitin tautologies. *Computational Complexity*, 20(4):649–678, 2011.
- [2] E. Amaldi and V. Kann. The complexity and approximability of finding maximum feasible subsystems of linear relations. *Theor. Comput. Sci.*, 147(1&2):181–210, 1995.
- [3] E. Amaldi and V. Kann. On the approximability of minimizing nonzero variables or unsatisfied relations in linear systems. *Theor. Comput. Sci.*, 209(1-2):237–260, 1998.

- [4] P. Austrin and S. Khot. A characterization of approximation resistance for even k -partite csps. In R. D. Kleinberg, editor, *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pages 187–196. ACM, 2013.
- [5] B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.
- [6] N. Creignou. A dichotomy theorem for maximum generalized satisfiability problems. *Journal of Computer and System Sciences*, 51(3):511–522, 1995.
- [7] A. De, E. Mossel, and J. Neeman. Majority is stablest: discrete and sos. In D. Boneh, T. Roughgarden, and J. Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 477–486. ACM, 2013.
- [8] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
- [9] K. M. Elbassioni, R. Raman, S. Ray, and R. Sitters. On the approximability of the maximum feasible subsystem problem with 0/1-coefficients. In C. Mathieu, editor, *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 1210–1219. SIAM, 2009.
- [10] V. Feldman, V. Guruswami, P. Raghavendra, and Y. Wu. Agnostic learning of monomials by halfspaces is hard. *SIAM J. Comput.*, 41(6):1558–1590, 2012.
- [11] S. Gaspers and S. Szeider. Kernels for global constraints. In T. Walsh, editor, *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 540–545. IJCAI/AAAI, 2011.
- [12] S. Gaspers and S. Szeider. Backdoors to acyclic SAT. In A. Czumaj, K. Mehlhorn, A. M. Pitts, and R. Wattenhofer, editors, *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*, volume 7391 of *Lecture Notes in Computer Science*, pages 363–374. Springer, 2012.
- [13] S. Gaspers and S. Szeider. Guarantees and limits of preprocessing in constraint satisfaction and reasoning. *Artif. Intell.*, 216:1–19, 2014.
- [14] M. Grohe. The structure of tractable constraint satisfaction problems. In R. Kralovic and P. Urzyczyn, editors, *MFCs 2006, Stará Lesná, Slovakia, August 28-September 1, 2006, Proceedings*, volume 4162 of *Lecture Notes in Computer Science*, pages 58–72. Springer, 2006.

- [15] F. Gurski and E. Wanke. The tree-width of clique-width bounded graphs without K_n , n . In U. Brandes and D. Wagner, editors, *Graph-Theoretic Concepts in Computer Science, 26th International Workshop, WG 2000, Konstanz, Germany, June 15-17, 2000, Proceedings*, volume 1928 of *Lecture Notes in Computer Science*, pages 196–205. Springer, 2000.
- [16] V. Guruswami and P. Raghavendra. Hardness of learning halfspaces with noise. *SIAM J. Comput.*, 39(2):742–765, 2009.
- [17] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [18] S. Khanna, M. Sudan, and D. P. Williamson. A complete classification of the approximability of maximization problems derived from boolean constraint satisfaction. In F. T. Leighton and P. W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 11–20. ACM, 1997.
- [19] S. Khot and R. Saket. Approximating csps using LP relaxation. In M. M. Halldórsson, K. Iwama, N. Kobayashi, and B. Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 822–833. Springer, 2015.
- [20] M. Lampis. Algorithmic meta-theorems for restrictions of treewidth. *Algorithmica*, 64(1):19–37, 2012.
- [21] H. W. Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of operations research*, 8(4):538–548, 1983.
- [22] D. Marx. Parameterized complexity and approximation algorithms. *Comput. J.*, 51(1):60–78, 2008.
- [23] S. Ordyniak, D. Paulusma, and S. Szeider. Satisfiability of acyclic and almost acyclic CNF formulas. *Theor. Comput. Sci.*, 481:85–99, 2013.
- [24] D. Paulusma, F. Slivovsky, and S. Szeider. Model counting for CNF formulas of bounded modular treewidth. In N. Portier and T. Wilke, editors, *STACS 2013, February 27 - March 2, 2013, Kiel, Germany*, volume 20 of *LIPICs*, pages 55–66. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [25] R. Pichler, S. Rümmele, S. Szeider, and S. Woltran. Tractable answer-set programming with weight constraints: bounded treewidth is not enough. *TPLP*, 14(2):141–164, 2014.
- [26] S. H. Sæther, J. A. Telle, and M. Vatshelle. Solving maxsat and #sat on structured CNF formulas. In C. Sinz and U. Egly, editors, *SAT 2014 - Vienna, Austria, July 14-17, 2014. Proceedings*, volume 8561 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2014.

- [27] M. Samer and S. Szeider. Constraint satisfaction with bounded treewidth revisited. *J. Comput. Syst. Sci.*, 76(2):103–114, 2010.
- [28] T. J. Schaefer. The complexity of satisfiability problems. In R. J. Lipton, W. A. Burkhard, W. J. Savitch, E. P. Friedman, and A. V. Aho, editors, *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, May 1-3, 1978, San Diego, California, USA*, pages 216–226. ACM, 1978.
- [29] F. Slivovsky and S. Szeider. Model counting for formulas of bounded clique-width. In L. Cai, S. Cheng, and T. W. Lam, editors, *ISAAC 2013, Hong Kong, China, Proceedings*, volume 8283 of *Lecture Notes in Computer Science*, pages 677–687. Springer, 2013.
- [30] S. Szeider. On fixed-parameter tractable parameterizations of SAT. In E. Giunchiglia and A. Tacchella, editors, *Theory and Applications of Satisfiability Testing, 6th International Conference, SAT 2003. Santa Margherita Ligure, Italy, May 5-8, 2003 Selected Revised Papers*, volume 2919 of *Lecture Notes in Computer Science*, pages 188–202. Springer, 2003.
- [31] S. Szeider. Not so easy problems for tree decomposable graphs. *CoRR*, abs/1107.1177, 2011.
- [32] S. Szeider. The parameterized complexity of constraint satisfaction and reasoning. In H. Tompits, S. Abreu, J. Oetsch, J. Pührer, D. Seipel, M. Umeda, and A. Wolf, editors, *INAP 2011, and WLP 2011, Vienna, Austria, September 28-30, 2011, Revised Selected Papers*, volume 7773 of *Lecture Notes in Computer Science*, pages 27–37. Springer, 2011.
- [33] S. Szeider. The parameterized complexity of k-flip local search for SAT and MAX SAT. *Discrete Optimization*, 8(1):139–145, 2011.
- [34] L. Trevisan. Inapproximability of combinatorial optimization problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 2004.